

## DETAILED ACTION

### *Remarks*

1. This communication is responsive to the amendment filed January 22<sup>nd</sup>, 2010. Claims 1, 2, 4-6, 9, 10, 12, 14-16, 19, 20, 22-27, 29, 30, 32-34, 37, 38, and 40-66 are pending. In the amendment filed January 22<sup>nd</sup>, 2010, Claims 1, 12, 22, 29, 38, 46, 51, 53, and 55 are amended, Claims 3, 7, 8, 11, 13, 17, 18, 21, 28, 31, 35, 36, and 39 are canceled, Claims 57-66 are new, and Claims 1, 12, 22, 29, 38, 46, 51, 53, and 55 are independent. This action is made FINAL.

### *Response to Arguments*

2. The Applicant's remarks filed January 22<sup>nd</sup>, 2010 with respect to Claims 1, 2, 4-6, 9, 10, 12, 14-16, 19, 20, 22-27, 29, 30, 32-34, 37, 38, and 40-66 have been fully considered but they are not persuasive. See below for a detailed response.

3. As to Applicant's arguments with respect to exemplary Claim 1 (including Claims 12, 29, 38, 51, and 53) for the prior art(s) allegedly not disclosing **"said at least one entry from said allocation layer is stored only on its corresponding computing device of a plurality of computing devices and said at least one entry from said B-link tree layer is replicated on each of said plurality of computing devices,"** the examiner respectfully disagrees. Lomet, col. 5, lines 38-44 was generally used to teach the allocation layer and B-link tree Layer (also termed "upper layer" in other claims). In the cited sections Lomet teaches about an RLOG that contains actions such as

updates/deletes/inserts (B-link tree/upper layer) as well as block allocates and frees (allocation layer). Thus, the RLOG of Lomet is seen as single log storing both a B-link tree/upper layer and an allocation layer. Each node in Lomet has its own RLOG (Lomet, col. 5, lines 52-61). Thus, every computing device has its own allocation layer stored thereon. Shaheen was brought into the rejections below to show that it is obvious for the B-link tree/upper layer to get copied to the other nodes. The cited sections of Shaheen (Shaheen, col. 4, lines 62-66 with Shaheen, col. 5, lines 12-14) teach that a log of data modifications are replicated across participating servers. The log in Shaheen that is getting replicated is only a log of data modifications (explicitly avoiding an allocation layer). It is obvious to include Shaheen into Lomet because it offers the obvious advantage of making Lomet distributed to increase the availability and reliability of distributed systems of Lomet (Shaheen, col. 1, lines 36-39). Also, replicating the allocation layer part of the log may be seen as wasting resources since the different nodes of Lomet may have different configurations (e.g. hard drive sizes) where block allocations/frees and sizes would be different between nodes and the allocs/frees would not correspond to every node. Thus, with Shaheen, the allocation layer is local to each machine that created the RLOG on its node and the B-link tree/upper layer is replicated to other servers/nodes as claimed.

4. As to Applicant's arguments with respect to exemplary Claim 22 for the prior art(s) allegedly not disclosing **"said entries from said allocation layer are stored only on their corresponding computing devices of a plurality of computing devices and said entries from said upper layer are replicated on each of said plurality of**

**computing devices,”** the examiner respectfully disagrees. This argument appears to have been met above.

5. As to Applicant's arguments with respect to exemplary Claim 46 for the prior art(s) allegedly not disclosing **“said entries from the upper layer are replicated to both said first server and said second server, while entries from said allocation layer at said first server are stored only on said first server, and entries from said allocation layer at said second server are stored only on said second server,”** the examiner respectfully disagrees. This argument appears to have been met above.

6. As to Applicant's arguments with respect to exemplary Claim 55 for the prior art(s) allegedly not disclosing **“said entries from the upper layer are replicated to both said first server location and said second server location , while entries from said allocation layer at said first server location are stored only on said first server location, and entries from said allocation layer at said second server location are stored only on said second server location,”** the examiner respectfully disagrees. This argument appears to have been met above.

7. The other claims argued merely because of a dependency on a previously argued claim(s) in the arguments presented to the examiner, filed January 22<sup>nd</sup>, 2010, are moot in view of the examiner's interpretation of the claims and art and are still considered rejected based on their respective rejections from at least a prior Office action (part(s) of recited again below).

***Specification***

8. The lengthy specification has not been checked to the extent necessary to determine the presence of all possible minor errors. Applicant's cooperation is requested in correcting any errors of which applicant may become aware in the specification.

***Claim Objections***

9. Claims 57-59 and 65 is objected to because of the following informalities:
- a. Claim 57 recites "pre-fix" while Claims 58 and 59 recite "prefix." It is unsure how these prefixes differ, if at all. Claim 57 or Claims 58 and 59 appear to contain a typographical error.
  - b. Claim 65 recites "The method according to claim1, wherein the B-link tree layer handles B-link tree a plurality of operations..." in lines 1-2. There appears to be a space missing after the word "claim" and an extraneous second "B-link tree" phrase. It appears the applicant intended to claim "The method according to claim\_1, wherein the B-link tree layer handles ~~B-link tree~~ a plurality of operations."

Appropriate correction is required.

***Claim Rejections - 35 USC § 112***

10. The following is a quotation of the first paragraph of 35 U.S.C. 112:

The specification shall contain a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the

art to which it pertains, or with which it is most nearly connected, to make and use the same and shall set forth the best mode contemplated by the inventor of carrying out his invention.

11. Claim 59 is rejected under 35 U.S.C. 112, first paragraph, as failing to comply with the written description requirement. The claim(s) contains subject matter which was not described in the specification in such a way as to reasonably convey to one skilled in the relevant art that the inventor(s), at the time the application was filed, had possession of the claimed invention. The examiner could not find support in the specification for a full prefix, let alone a full prefix indicative that the single B-link tree operation has been logically completed. Paragraph [0083] spanning pages 24 and 25 of the specification appeared to be the only section dealing with prefixes.

### ***Claim Rejections - 35 USC § 103***

12. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

13. This application currently names joint inventors. In considering patentability of the claims under 35 U.S.C. 103(a), the examiner presumes that the subject matter of the various claims was commonly owned at the time any inventions covered therein were made absent any evidence to the contrary. Applicant is advised of the obligation under 37 CFR 1.56 to point out the inventor and invention dates of each claim that was not commonly owned at the time a later invention was made in order for the examiner to

consider the applicability of 35 U.S.C. 103(c) and potential 35 U.S.C. 102(e), (f) or (g) prior art under 35 U.S.C. 103(a).

14. Claims 1, 2, 4-6, 9, 10, 12, 14-16, 19, 20, 29, 30, 32-34, 37, 38, 40-45, 51-54, 57-62, and 64-66 are rejected under 35 U.S.C. 103(a) as being unpatentable over U.S. Patent No. 5,485,608 (Lomet et al.) in view of "Efficient Locking for Concurrent Operations on B-Trees" (Lehman et al.), further in view of U.S. Patent No. 5,434,994 (Shaheen et al.).

For **Claim 1**, Lomet teaches: "A method for logging while updating...via a plurality of data transactions, [Lomet, col. 7, lines 26-34, write ahead logging – log written before blocks during update process] the method comprising:

- generating a plurality of log entries corresponding to a plurality of data transactions, the data transactions to be carried out on a plurality of nodes of a B-link tree, wherein the data transactions are associated with a single B-link tree operation on said B-link tree, [Lomet, col. 1, lines 29-48 with Lomet, col. 22, lines 25-30 with Lomet, col. 19, lines 35-37, (B-link tree taught by lehman below) log records created when updating (e.g. RLOG)] wherein said plurality of log entries include at least one entry from an allocation layer [Lomet, col. 1, lines 65-67 with Lomet, col. 5, lines 38-44, RLOG data entry for block alloc's/frees] and at least one entry from a B-link tree layer, [Lomet, col. 19, lines 35-37 with Lomet, col. 5, lines 38-44, action (e.g. update, delete) done to Lehman's b-link tree] wherein said at least one entry from said allocation layer is stored only on its

corresponding computing device of a plurality of computing devices [Lomet, col. 2, lines 50-53 with Lomet, col. 5, lines 51-62, RLOG on each node (where only modifications (to records) are replicated according to Shaheen below)]

- ...storing said plurality of log entries in a single log, the single log being a partially persistent log, wherein a tail of said partially persistent log is in a memory, and said partially persistent log comprises of said memory and a persistent storage; [Lomet, col. 4, lines 28-30 with Lomet, col. 4, lines 64-67 with Lomet, col. 5, lines 1-3 with Lomet, col. 12, lines 20-23 with Lomet, col. 7, lines 26-34 with Lomet, col. 19, lines 35-37, RLOG of entries is singular to each node where its tail is buffered (e.g. volatile), then the RLOG is flushed to one persistent shared RLOG disk when associated transaction commits]
- associating said plurality of log entries with each other for determining whether the single B-link tree operation has been completed; [Lomet, col. 1, lines 29-48 with Lomet, col. 22, lines 25-30 with Lomet, col. 19, lines 35-37, logged changes associated with a transaction and committing guarantees that the effects of the transaction are persistent] and
- wherein a boundary between said memory and said persistent storage in said partially persistent log varies over time" [Lomet, col. 4, lines 28-30 with Lomet, col. 4, lines 64-67 with Lomet, col. 5, lines 1-3 with Lomet, col. 12, lines 20-23 with Lomet, col. 19, lines 35-37 with Lomet, col. 16, lines 23-25 see also Lomet, col. 19, lines 45-51 with Lomet, col. 20, lines 20-26, RLOG buffer occasionally (when commits) flushes entries to persistent storage. This can be viewed as a

boundary varying over time (as new transactions start and others commit over time by database usage)).

Lomet discloses the above limitations but does not expressly teach:

- "...a B-link tree
- ...and said at least one entry from said B-link tree is replicated on each of said plurality of computing devices."

With respect to Claim 1, an anALOGous art, Lehman, teaches:

- "...a B-link tree" [Lehman, p. 657, section 3.3, b-trees (specifically including b-link trees) are favorable for database applications].

With respect to Claim 1, an anALOGous art, Shaheen, teaches:

- "...and said at least one entry from said B-link tree is replicated on each of said plurality of computing devices" [Shaheen, col. 4, lines 62-66 with Shaheen, col. 5, lines 12-14, log of data modifications are replicated across participating servers].

It would have been obvious to one of ordinary skill in the art at the time of invention to combine Lehman and Shaheen with Lomet because both inventions are directed towards storing information in a database.

Lehman's and Shaheen's inventions would have been expected to successfully work well with Lomet's invention because the inventions use databases. Lomet discloses methods and an apparatus for updating information in a computer system using logs and state identifiers comprising RLOGs, ULOGs and ALOGs, however Lomet does not expressly disclose the use of a B-link-tree to store the data in the database of Lomet nor replicating the modifications to the other computers. Lehman discloses



efficient locking for concurrent operations on B-trees comprising a B-link-tree. Shaheen discloses a system and method for maintaining replicated data coherency in a data processing system comprising replicating data and logs across multiple servers.

It would have been obvious to one of ordinary skill in the art at the time of invention to take the b-link-tree from Lehman and install it into the invention of Lomet, thereby offering the obvious advantage of a guaranteed small (average) search, insertion, and deletion time for the database of Lomet (see front page of Lehman).

It would have been obvious to one of ordinary skill in the art at the time of invention to take the replication of data and logs across multiple servers from Shaheen and install it into the invention of Lomet, thereby offering the obvious advantage of making Lomet distributed to increase the availability and reliability of distributed systems of Lomet (Shaheen, col. 1, lines 36-39).

**Claim 2** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 1, further including periodically truncating the log" [Lomet, cols. 11-12, lines 60-12].

**Claim 4** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 1, further including discarding one of the log entries from the log when the corresponding data transaction has been carried out on said B-link tree" [Lomet, col. 14, lines 9-12].

**Claim 5** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 1, wherein said storing includes storing said log

entries into the log before the corresponding data transactions are carried out on said B-link tree" [Lomet, col. 7, lines 26-34].

**Claim 6** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 1, further including caching data of said data transactions before said data transactions are carried out on said B-link tree" [Lomet, col. 7, lines 26-34 with Lomet, col. 5, lines 1-11].

**Claim 9** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 1, further including maintaining a log sequence number with each of said plurality of log entries, uniquely identifying one of the log entries" [Lomet, col. 5, lines 44-52].

**Claim 10** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 1, wherein said method is embodied in instructions that are stored on a computer readable medium in physical memory" [Lomet, col. 4, lines 17-30, memory for the corresponding processor to execute routines].

For **Claim 12**, Lomet teaches: "A method for logging while updating...via a plurality of data transactions, [Lomet, col. 7, lines 26-34] the method comprising:

- generating a plurality of log entries corresponding to a plurality of data transactions, the data transactions to be carried out on a plurality of nodes of a B-link tree, wherein the data transactions are associated with a single B-link operation on said B-link tree, [Lomet, col. 1, lines 29-48 with Lomet, col. 22, lines 25-30 with Lomet, col. 19, lines 35-37 with Lomet, col. 19, lines 45-51] wherein

said plurality of log entries include at least one entry from an allocation layer [Lomet, col. 1, lines 65-67 with Lomet, col. 5, lines 38-44] and at least one entry from a B-link tree layer, [Lomet, col. 19, lines 35-37 with Lomet, col. 5, lines 38-44] wherein said at least one entry from said allocation layer is stored only on its corresponding computing device of a plurality of computing devices [Lomet, col. 2, lines 50-53 with Lomet, col. 5, lines 51-62]

- ...storing said plurality of log entries in a single log, the single log being a partially persistent log, wherein a tail of said partially persistent log is in a memory, and said partially persistent log comprises of said memory and a persistent storage; [Lomet, col. 4, lines 28-30 with Lomet, col. 4, lines 64-67 with Lomet, col. 5, lines 1-3 with Lomet, col. 12, lines 20-23 with Lomet, col. 7, lines 26-34 with Lomet, col. 19, lines 35-37]
- associating said plurality of log entries with each other for determining whether the single B-link operation has been completed; [Lomet, col. 1, lines 29-48 with Lomet, col. 22, lines 25-30 with Lomet, col. 19, lines 35-37 with Lomet, col. 19, lines 45-51] and
- wherein a boundary between said memory and said persistent storage in said partially persistent log varies over time [Lomet, col. 4, lines 28-30 with Lomet, col. 4, lines 64-67 with Lomet, col. 5, lines 1-3 with Lomet, col. 12, lines 20-23 with Lomet, col. 19, lines 35-37 with Lomet, col. 16, lines 23-25 see also Lomet, col. 19, lines 45-51 with Lomet, col. 20, lines 20-26]

- periodically flushing data corresponding to data transactions represented by the finite log to persistent storage; [Lomet, col. 14, lines 9-17] and
- truncating said finite log in coordination with said flushing" [Lomet, col. 14, lines 9-17].

Lomet discloses the above limitations but does not expressly teach:

- "...a B-link tree
- ... and said at least one entry from said B-link tree is replicated on each of said plurality of computing devices."

With respect to Claim 12, an anALOGous art, Lehman, teaches:

- "...a B-link tree" [Lehman, p. 657, section 3.3].

With respect to Claim 12, an anALOGous art, Shaheen, teaches:

- "...and said at least one entry from said B-link tree is replicated on each of said plurality of computing devices" [Shaheen, col. 4, lines 62-66 with Shaheen, col. 5, lines 12-14, log of data modifications are replicated across participating servers].

It would have been obvious to one of ordinary skill in the art at the time of invention to combine Lehman and Shaheen with Lomet because both inventions are directed towards storing information in a database.

Lehman's and Shaheen's inventions would have been expected to successfully work well with Lomet's invention because the inventions use databases. Lomet discloses methods and an apparatus for updating information in a computer system using logs and state identifiers comprising RLOGs, ULOGs and ALOGs, however Lomet does not expressly disclose the use of a B-link-tree to store the data in the database of

Lomet nor replicating the modifications to the other computers. Lehman discloses efficient locking for concurrent operations on B-trees comprising a B-link-tree. Shaheen discloses a system and method for maintaining replicated data coherency in a data processing system comprising replicating data and logs across multiple servers.

It would have been obvious to one of ordinary skill in the art at the time of invention to take the b-link-tree from Lehman and install it into the invention of Lomet, thereby offering the obvious advantage of a guaranteed small (average) search, insertion, and deletion time for the database of Lomet (see front page of Lehman).

It would have been obvious to one of ordinary skill in the art at the time of invention to take the replication of data and logs across multiple servers from Shaheen and install it into the invention of Lomet, thereby offering the obvious advantage of making Lomet distributed to increase the availability and reliability of distributed systems of Lomet (Shaheen, col. 1, lines 36-39).

**Claim 14** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 12, further including discarding one of the log entries from the finite log when the corresponding data transaction has been carried out on said B-link tree" [Lomet, col. 14, lines 9-17].

**Claim 15** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 12, wherein said storing includes storing said log entries into the finite log before the corresponding data transactions are carried out on said B-link tree" [Lomet, col. 7, lines 26-34].

**Claim 16** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 12, further including caching data of said data transactions before said data transactions are carried out on said B-link tree" [Lomet, col. 7, lines 26-34 with Lomet, col. 5, lines 1-11].

**Claim 19** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 12, further including maintaining a log sequence number with each of said plurality of log entries, uniquely identifying one of the log entries" [Lomet, col. 5, lines 44-52].

**Claim 20** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 12, wherein said method is embodied in instructions that are stored on a computer readable medium in physical memory" [Lomet, col. 4, lines 17-30, memory for the corresponding processor to execute routines].

**Claims 29, 30, 32-34, and 37** encompass substantially the same scope of the invention as that of Claims 1, 2, 4-6, and 9, respectfully, in addition to a server and some actions for performing the method steps of Claims 1, 2, 4-6, and 9, respectfully. Therefore, Claims 29, 30, 32-34, and 37 are rejected for the same reasons as stated above with respect to Claims 1, 2, 4-6, and 9, respectfully.

**Claims 38 and 40-42** encompass substantially the same scope of the invention as that of Claims 12 and 14-16, respectfully, in addition to a server and some objects for performing the method steps of Claims 12 and 14-16, respectfully. Therefore, Claims 38

and 40-42 are rejected for the same reasons as stated above with respect to Claims 12 and 14-16, respectfully.

**Claim 43** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A server according to claim 38, further including storing said plurality of log entries in an intermediate memory previous to storing said plurality of log entries in the finite storage log" [Lomet, col. 7, lines 26-34 with Lomet, col. 19, lines 35-37].

**Claim 44** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A server according to claim 43, wherein said plurality of log entries are moved from intermediate memory to the finite storage log after the data transactions commit" [Lomet, col. 16, lines 23-25 see also Lomet, col. 19, lines 45-51 with Lomet, col. 20, lines 20-26].

**Claim 45** encompasses substantially the same scope of the invention as that of Claim 19, in addition to a server and some objects for performing the method steps of Claim 19. Therefore, Claim 45 is rejected for the same reasons as stated above with respect to Claim 19.

**Claims 51 and 52** encompass substantially the same scope of the invention as that of Claims 1 and 2, respectfully, in addition to a computing device and some means for performing the method steps of Claims 1 and 2, respectfully. Therefore, Claims 51 and 52 are rejected for the same reasons as stated above with respect to Claims 1 and 2, respectfully.

**Claims 53 and 54** encompass substantially the same scope of the invention as that of Claims 12 and 14, respectfully, in addition to a computing device and some

means for performing the method steps of Claims 12 and 14, respectfully. Therefore, Claims 53 and 54 are rejected for the same reasons as stated above with respect to Claims 12 and 14, respectfully.

**Claim 57** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 1, wherein said plurality of log entries are associated with each other via a pre-fix" [Lomet, Fig. 3, TID or LSN]

**Claim 58** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 57, wherein a partial prefix in the single log is indicative that the single B-link tree operation has not been completed" [Lomet, Fig. 3 with Lomet, col. 19, lines 20-27, all transactions prefixed with start\_transaction log record. This is indicative that the operation has not been completed since it is on the RLOG].

**Claim 59** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 57, wherein a full prefix in the single log is indicative that the single B-link tree operation has been logically completed" [Lomet, Fig. 5 with Lomet, col. 19, lines 20-27 with Lomet, col. 20, lines 15-26, prepare log record prefixes the commit log record this indicates that the transaction has *logically* committed].

**Claim 60** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 1, wherein each of said plurality of log entries corresponds to an atomic operation" [Shaheen, col. 7, lines 49-54 updates in transaction are guaranteed atomicity. Also, Lomet suggests atomicity at col. 1, lines 29-48].



**Claim 61** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 1, wherein data of said data transactions are cached in a cache memory before said data transactions are carried out on said B-link tree" [Lomet, col. 7, lines 26-34 with Lomet, col. 5, lines 1-11].

**Claim 62** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 1, wherein at least one entry from the allocation layer is written to a disk on its corresponding computing device when the single B-link tree operation becomes persistent" [Lomet, col. 5, lines 37-44 with Lomet, col. 7, lines 27-33 with Lomet, col. 19, lines 45-51 with Lomet, col. 20, lines 21-26, RLOG, containing allocation layer of log, is written to buffer prior to commit (write ahead logging), then on commit, the RLOG is stored to disk].

**Claim 64** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 1, wherein said storing further comprising:

- storing a first log entry of said plurality of log entries in the tail of said partially persistent log when its respective data transaction has not been persisted; [Lomet, col. 19, lines 35-40 with Lomet, col. 18, lines 1-12 with Lomet, col. 14, lines 3-10 with Lomet, col. 5, lines 45-51, LSN inherent from record position where newer records are at the tail end of the RLOG buffer] and
- storing a second log entry of said plurality of log entries in the persistent storage when its respective data transaction has been persisted" [Lomet, col. 5, lines 37-44 with Lomet, col. 7, lines 27-33 with Lomet, col. 19, lines 45-51 with Lomet, col.

20, lines 21-26, RLOG is written to buffer prior to commit (write ahead logging), then on commit, the RLOG is stored to disk].

**Claim 65** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim1, wherein the B-link tree layer handles B-link tree a plurality of operations comprising an insert node operation, a lookup node operation, and a delete node operation" [Lomet, col. 5, lines 37-44, inserts, deletes. Additionally, in order to delete a record it must be located first (look up). Lehman also teaches a lookup/search operation on p. 659, section 4].

**Claim 66** can be mapped to Lomet (as modified by Lehman and Shaheen) as follows: "A method according to claim 1, wherein the allocation layer handles a plurality of operations comprising an allocate disk space operation, a deallocate disk space operation, a read from the allocated disk space operation and a write to the allocated disk space operation" [Lomet, col. 5, lines 31-38, allocate/free. Additionally, freeing/allocating requires writing and reading].

15. Claims 22, 23, 26, 27, 46, 47, 50, 55 and 56 are rejected under 35 U.S.C. 103(a) as being unpatentable over U.S. Patent No. 5,485,608 (Lomet et al.) in view of U.S. Patent No. 5,434,994 (Shaheen et al.).

For **Claim 22**, Lomet teaches: "A method for logging while updating a data structure via a plurality of data transactions, [Lomet, col. 7, lines 26-34, write ahead logging for transactions] the method comprising:

- ...generating a plurality of log entries corresponding to a plurality of data transactions, the data transactions to be carried out on said data structure, wherein the data transactions are associated with a single operation on the data structure; [Lomet, col. 1, lines 29-48 with Lomet, col. 22, lines 25-30 with Lomet, col. 19, lines 35-37 with Lomet, col. 19, lines 45-51, transaction being a series of changes committed to a database (has a structure). Committing guarantees that the effects of the transaction are persistent. Each change is logged in at least an RLOG] and
- ...where the single log is partitioned into an upper layer and an allocation layer, [Lomet, col. 5, lines 31-45 with Lomet, Fig. 3, data updates/inserts/deletes (modifications) and block allocates/frees in RLOG] ... wherein the single log includes log entries from both the upper layer and allocation layer [Lomet, col. 5, lines 15-22 with Lomet, Fig. 3, both in RLOG (REDO\_DATA)]...wherein said entries from said allocation layer are stored only on their corresponding computing device of a plurality of computing devices [Lomet, col. 2, lines 50-53 with Lomet, col. 5, lines 51-62, RLOG on each node (where only modifications (to records) are replicated according to Shaheen below)]
- ...associating said plurality of log entries with each other for determining whether the single operation on the data structure has been completed; [Lomet, col. 1, lines 29-48 with Lomet, col. 22, lines 25-30 with Lomet, col. 19, lines 35-37 with Lomet, col. 19, lines 45-51 with Lomet, col. 20, lines 21-26, logged changes

associated with a transaction and committing guarantees that the effects of the transaction are persistent] and

- wherein said single log is a partially persistent log, where a tail of said partially persistent log is in a memory, and said partially persistent log comprises of said memory and a persistent storage" [Lomet, col. 4, lines 28-30 with Lomet, col. 4, lines 64-67 with Lomet, col. 5, lines 1-3 with Lomet, col. 12, lines 20-23 with Lomet, col. 7, lines 26-34 with Lomet, col. 19, lines 35-37, RLOG is of a buffer (memory). It is partially persistent in that it persists until the memory is no longer powered or it moves to the ALOG. The RLOG is flushed to a disk upon a commit instruction (incase it needs to be recovered it is made persistent)].

Lomet discloses the above limitations but does not expressly teach:

- "...replicating updates to the data structure to a first server location and a second server location;
- maintaining a single log...at each of said first and second server locations...and said entries from said upper layer are replicated on each of said plurality of computing devices."

With respect to Claim 22, an anALOGous art, Shaheen, teaches:

- "...replicating updates to the data structure to a first server location and a second server location; [Shaheen, col. 7, lines 50-55, each server updates from the merged log]
- maintaining a single log...at each of said first and second server locations [Shaheen, col. 4, lines 62-66, modification logs are merged and sent to each

participating server]... and said entries from said upper layer are replicated on each of said plurality of computing devices" [Shaheen, col. 4, lines 62-66 with Shaheen, col. 5, lines 12-14, log of data modifications are replicated across participating servers].

It would have been obvious to one of ordinary skill in the art at the time of invention to combine Shaheen with Lomet because both inventions are directed towards recovering data upon failure.

Shaheen's invention would have been expected to successfully work well with Lomet's invention because both inventions use computers using logs. Lomet discloses methods and an apparatus for updating information in a computer system using logs and state identifiers comprising RLOGs, ULOGs and ALOGs, however Lomet does not expressly disclose replicating data on multiple servers, or maintaining a single log on those servers (although, Lomet is seen as teaching the single log, Shaheen makes it for multiple servers). Shaheen discloses a system and method for maintaining replicated data coherency in a data processing system comprising replicating data and logs across multiple servers.

It would have been obvious to one of ordinary skill in the art at the time of invention to take the replication of data and logs across multiple servers from Shaheen and install it into the invention of Lomet, thereby offering the obvious advantage of making Lomet distributed to increase the availability and reliability of distributed systems of Lomet (Shaheen, col. 1, lines 36-39).

**Claim 23** can be mapped to Lomet (as modified by Shaheen) as follows: "A method according to claim 22, further including recovering the data structure after a failure by performing parallel recovery operations by each of said first and second server locations" [Lomet, col. 21, lines 35-39 with Shaheen, col. 4, lines 60-66 with Shaheen, col. 7, lines 49-57].

**Claim 26** can be mapped to Lomet (as modified by Shaheen) as follows: "A method according to claim 22, wherein the allocation layer handles at least one of (A) an allocate disk space operation, (B) a deallocate disk space operation, (C) a read from the allocated disk space operation and (D) a write to the allocated disk space operation" [Lomet, col. 5, lines 31-38].

**Claim 27** can be mapped to Lomet (as modified by Lehman) as follows: "A method according to claim 22, wherein said method is embodied in instructions that are stored on a computer readable medium in physical memory" [Lomet, col. 4, lines 17-30, memory for the corresponding processor to execute routines].

**Claims 46, 47, and 50** encompass substantially the same scope of the invention as that of Claims 22, 23, and 26, respectively, in addition to a server and some objects for performing the method steps of Claims 22, 23, and 26, respectively. Therefore, Claims 46, 47, and 50 are rejected for the same reasons as stated above with respect to Claims 22, 23, and 26, respectively. Additionally, Claim 46 recites "wherein said entries from the upper layer are replicated to both said first server and said second server, [Shaheen, col. 4, lines 62-66 with Shaheen, col. 5, lines 12-14, log of data modifications are replicated across participating servers] while entries from said allocation layer at

said first server are stored only on said first server, and entries from said allocation layer at said second server are stored only on said second server" [Lomet, col. 2, lines 50-53 with Lomet, col. 5, lines 51-62, RLOG on each node (where only modifications (to records) are replicated according to Shaheen above)] which is taught as shown above (similar to the above claims). Also, Claim 46 has an additional limitation not found in Claim 22, but is also taught by Lomet : "wherein a boundary between said memory and said persistent storage in said partially persistent log varies over time" [Lomet, col. 4, lines 28-30 with Lomet, col. 4, lines 64-67 with Lomet, col. 5, lines 1-3 with Lomet, col. 12, lines 20-23 with Lomet, col. 19, lines 35-37 with Lomet, col. 16, lines 23-25 see also Lomet, col. 19, lines 45-51 with Lomet, col. 20, lines 20-26].

**Claims 55 and 56** encompass substantially the same scope of the invention as that of Claims 22 and 23, respectfully, in addition to a computing device and some means for performing the method steps of Claims 22 and 23, respectfully. Therefore, Claims 55 and 56 are rejected for the same reasons as stated above with respect to Claims 22 and 23, respectfully. Additionally, Claim 55 recites "wherein said entries from the upper layer are replicated to both said first server location and said second server location, [Shaheen, col. 4, lines 62-66 with Shaheen, col. 5, lines 12-14, log of data modifications are replicated across participating servers] while entries from said allocation layer at said first server location are stored only on said first server location, and entries from said allocation layer at said second server location are stored only on said second server location" [Lomet, col. 2, lines 50-53 with Lomet, col. 5, lines 51-62, RLOG on each node (where only modifications (to records) are replicated according to

Shaheen above)] which is taught as shown above (similar to the above claims). Also, Claim 55 recites “wherein said single log is a partially persistent log that has a boundary that changes over time between a persistent and non-persistent memory” which is taught by [Lomet, col. 4, lines 28-30 with Lomet, col. 4, lines 64-67 with Lomet, col. 5, lines 1-3 with Lomet, col. 12, lines 20-23 with Lomet, col. 19, lines 35-37 with Lomet, col. 16, lines 23-25 see also Lomet, col. 19, lines 45-51 with Lomet, col. 20, lines 20-26] (similar to the above claims).

16. Claims 24, 25, 48, and 49 are rejected under 35 U.S.C. 103(a) as being unpatentable over U.S. Patent No. 5,485,608 (Lomet et al.) in view of U.S. Patent No. 5,434,994 (Shaheen et al.), further in view of Efficient Locking for Concurrent Operations on B-Trees (Lehman et al.).

For **Claim 24**, Lomet (as modified by Shaheen) teaches: “A method according to claim 22.”

Lomet (as modified by Shaheen) discloses the above limitation but does not expressly teach: “wherein said data structure is a B-link tree.”

With respect to Claim 24, an anALOGous art, Lehman, teaches: “wherein said data structure is a B-link tree” [Lehman, p. 657, section 3.3].

It would have been obvious to one of ordinary skill in the art at the time of invention to combine Lehman with Lomet (as modified by Shaheen) because both inventions are directed towards storing information in a database.



Lehman's invention would have been expected to successfully work well with Lomet (as modified by Shaheen)'s invention because both inventions use databases. Lomet (as modified by Shaheen) discloses methods and an apparatus for updating information in a computer system using logs and state identifiers comprising RLOGs, ULOGs and ALOGs, however Lomet (as modified by Shaheen) does not expressly disclose the use of a B-link-tree to store the data in the database of Lomet (as modified by Shaheen). Lehman discloses efficient locking for concurrent operations on B-trees comprising a B-link-tree.

It would have been obvious to one of ordinary skill in the art at the time of invention to take the b-link-tree from Lehman and install it into the invention of Lomet (as modified by Shaheen), thereby offering the obvious advantage of a guaranteed small (average) search, insertion, and deletion time for the database of Lomet (as modified by Shaheen).

**Claim 25** can be mapped to Lomet (as modified by Shaheen and Lehman) as follows: "A method according to claim 24, wherein the upper layer is a B-link tree layer that handles B-link tree operations" [Lomet, col. 5, lines 39-44].

**Claims 48 and 49** encompass substantially the same scope of the invention as that of Claims 24 and 25, respectfully, in addition to a server and some objects for performing the method steps of Claims 24 and 25, respectfully. Therefore, Claims 48 and 49 are rejected for the same reasons as stated above with respect to Claims 24 and 25, respectfully.

17. Claim 63 is rejected under 35 U.S.C. 103(a) as being unpatentable over U.S. Patent No. 5,485,608 (Lomet et al.) in view of "Efficient Locking for Concurrent Operations on B-Trees" (Lehman et al.), in view of U.S. Patent No. 5,434,994 (Shaheen et al.), further in view of U.S. Patent No. 6,792,436 (Zhu et al.).

For **Claim 63**, Lomet (as modified by Lehman and Shaheen) teaches: "A method according to claim 1, wherein the data transaction comprising an update, [Lomet, col. 5, lines 40-45] the method further comprising:

- ...storing a log entry corresponding to the update in the tail of said partially persistent log; [Lomet, col. 19, lines 35-40 with Lomet, col. 18, lines 1-12 with Lomet, col. 14, lines 3-10 with Lomet, col. 5, lines 45-51, LSN inherent from record position where newer records are at the tail end of the RLOG] and
- writing the log entry corresponding to the update in the persistent storage immediately before the updated cached data block is persisted" [Lomet, col. 5, lines 37-44 with Lomet, col. 7, lines 27-33 with Lomet, col. 19, lines 45-51 with Lomet, col. 20, lines 21-26, RLOG, containing allocation layer of log, is written to buffer prior to commit (write ahead logging), then on commit, the RLOG is stored to disk].

Lomet (as modified by Lehman and Shaheen) discloses the above limitations but does not explicitly teach:

- "...caching the data block in a cache memory;
- applying the update to the cached data block in the cache memory."

With respect to Claim 63, an anALOGous art, Zhu, teaches:

- "...caching the data block in a cache memory; [Zhu, col. 5, lines 42-50 with Zhu, col. 6, lines 41-43 with Zhu, col. 7, lines 15-21, cached object in transactional cache means it was cached]
- applying the update to the cached data block in the cache memory" [Zhu, col. 5, lines 42-50 with Zhu, col. 6, lines 41-43 with Zhu, col. 7, lines 15-21, transactional cache isolates changes to cached objects until committed].

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Zhu and Lomet (as modified by Lehman and Shaheen) before him/her to combine Zhu with Lomet (as modified by Lehman and Shaheen) because the inventions are directed towards database processing.

Zhu's invention would have been expected to successfully work well with Lomet (as modified by Lehman and Shaheen)'s invention because the inventions use databases on a network. Lomet (as modified by Lehman and Shaheen) discloses a methods and an apparatus for updating information in a computer system using logs and state identifiers comprising RLOGs, ULOGs and ALOGs. However, Lomet (as modified by Lehman and Shaheen) does not expressly disclose a cache for data block updates. Zhu discloses a method for synchronizing multiple software caches in a memory (title) comprising a transactional cache for database data to isolate local user changes until commit.

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Zhu and Lomet (as modified by Lehman and Shaheen) before him/her to take the transactional caches and updating from Zhu and

install them into the teachings of Lomet (as modified by Lehman and Shaheen), thereby offering the obvious advantage of protecting data integrity between users of the databases.

18. Applicant's amendment necessitated the new ground(s) of rejection presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire **THREE MONTHS** from the mailing date of this action. In the event a first reply is filed within **TWO MONTHS** of the mailing date of this final action and the advisory action is not mailed until after the end of the **THREE-MONTH** shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than **SIX MONTHS** from the date of this final action.

**Conclusion**

19. Any prior art made of record and not relied upon is considered pertinent to applicant's disclosure. Applicant is advised that, although not used in the rejections above, prior art cited on any PTO-892 form and not relied upon is considered materially relevant to the applicant's claimed invention and/or portions of the claimed invention.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to BRENT STACE whose telephone number is (571)272-8372 and fax number is 571-273-8372. The examiner can normally be reached on M-F 9am-5:30pm EST.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Apu M. Mofiz can be reached on 571-272-4080. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/BRENT STACE/ (4/8/10)  
Examiner, Art Unit 2161  
Granted Temp. FSA